NARDIS (1980)

DOUGLAS LUMAN

Bill Evans, live at the Keystone Korner, San Francisco
August 31–September 8, 1980

Works Statement

This sequence of "broken" (we are doing a jazz, after all) triolets is derived from a process using OpenAI's Whisper speech-to-text model[1] run on the last six versions of the song "Nardis" performed by jazz pianist Bill Evans with his last trio at the Keystone Korner (San Francisco) one week before his death in September 1980. "Nardis" is one of the songs most closely associated with Evans, even though Miles Davis wrote the chart in 1958 for the saxophonist Cannonball Adderley. The version of "Nardis" played on September 8, 1980, would be the last recording that Evans would make before his death.

I developed an interest in this song as an introspective look into Evans's awareness of his approaching death, a text-less entry into the artistic tradition of the *memento mori*. Only after sitting with my thoughts about these performances and repeatedly listening to these (and other) versions of it did I discover others' writing—particularly Jeff Meshel's in-depth treatment[2] of several studio and live recordings. We both heard an exploration of a man confronting a force much larger than just a set of late-career recording dates.

Given that these recordings contain no vocal parts (save for a brief exchange with the audience after the first night of the stand), transcribing text is impossible. However, it is intrinsic to my practice to push technology to places

1. https://platform.openai.com/docs/models/whisper-1
2. https://jmeshel.com/124-bill-evans-nardis/

where it functions, as I tell my students often enough, not as *intended* but as *designed*. The gap exploited here is the ability to instruct the OpenAI model to engage in wild speculation, turning sounds into linguistic problems for the model to solve, interpret, and return.

Of the various parameters (i.e., knobs, dials) that one can turn is the *temperature* of the model results: the vehicle allowing users to force the model to engage in the kind of guessing that can generate linguistic data in the music where none exists. Using the OpenAI Whisper API[3,4] via short Python program, I divided audio files into 10-minute segments and "played" segments of the MP3 files to the API with the temperature parameter changed to its maximum setting, 2.0 on a scale of 0.0–2.0. This value causes the model to take risks in interpretation far beyond recognizing the audio as music, which was Whisper's standard response in the beginning. While other values yielded some results, none were as useful as those created when the *temperature* was turned all the way up.

As the API is often used to transcribe standard audio fare (meetings, interviews, etc.), users are also allowed to submit *prompts* with each request. Often, these contain context about what the audio represents and any additional acronyms or information that would assist the model in making an accurate transcription. Being direct with Whisper (i.e., writing that "This is the jazz song 'Nardis,' played by Bill Evans") created less-valuable results, and the project required trying several different kinds of prompts to create more-or-less reliable (if not predictable) outcomes.

This sequence was composed over the course of a year and a half of trial and error, abandonment, and the traditional de/re/prioritization of creative projects. In that kind of geological time, models often change and approaches that once worked well commonly decrease in efficacy. While my *prompts* early on were much more abstract ("This recording features the pianist Bill Evans being interviewed . . . ,""what do you see,""how would you

3. API: Application Programming Interface. APIs constitute how a programmer might "speak to" a service such as OpenAI from other software using a commonly-accepted, public set of methods.

4. https://platform.openai.com/docs/api-reference/audio

describe where you are . . ."), the solution resulting in generating the most text came from reading OpenAI documentation about results including filler words ("like,""um," etc.), which resulted in the most generative prompt for these files: the anti-climactic "Um, let me think."

The Python program written to send and receive messages from Whisper printed the resulting model guesses to the screen, which were often a mixture of several languages and unrelated words. I followed my traditional method of copying these words to my writing notebook to form a lexicon, file by file. I ran the program many times per audio file until I felt that my collection of language was sufficient to begin to find patterns between writing sessions. In the medium run, patterns emerged that allowed me to create the text of the work. Due to the repetition of certain ideas, words, or phrases, I chose an equally repetitive form, the triolet, which allowed me to further narrow my decision space to create the sequence presented here.

* * *

August 31, 1980

What kind of meandering should a flock do.
If anything, it's your internal dog that makes you wonder
how many shepherds govern any given acre of emergence,
what kind of meandering you're supposed to do.
It's a voyage. It's underway. How do you prepare for it
except pray. How do you start praising a thing
with a vocabulary that meanders like a flock does, or,
like any dog, gets to wandering.

September 1, 1980

You whistle a basic six note phrase. It
becomes your anthem. Is this what
prayer looks like. Whistle a basic six note phrase—
your anthem—suddenly, a decent moon appears
and there's a film of light over everything.
Each hour, afraid of the previous one, moves
on only to find it again. You whistle a little phrase.
It becomes your anthem. This is what a prayer looks like.

September 2, 1980

If you improvise another allegory for the audience
all they do is clap. Somebody has to be the fool,
but I am too tired to grab at the robes of saints.
So, I improvise another allegory for the audience,
shoving a few more variables rendered in the spirit's
vocabulary or whatever language causes anyone to listen
to another improvised allegory for an audience.
All they do is clap.

September 6, 1980

If you listen to any gods, it ends
up like this: struggling with a manuscript.
Like any good guidebook—
if you listen to any gods—
the answers are always in the back.
Do an exercise: turn off the lights. Describe
the sound. Reproduce it faithfully. Are any gods
listening. How does the exercise end.

September 7, 1980

Today I have to wake up because
tomorrow I have to wake up. I presume
the geometry of a clock. So, I
wake up because I expect to be greeted
by either a circle or some stones. Now I'm out here
with my palms up. How can I be the only
one afraid of piety. I have to
wake up because tomorrow.

September 8, 1980

Question: what is a silhouette, but an illusion
taken at a slow tempo. How palpable it feels
when time—an empty vessel—no longer fills itself.
You question that silhouette. You speak
at a slow tempo. You ask how *urgently* and *gently* are
synonyms. Their music is as familiar as
your own shadow. Ask: is there a word for this
slow tempo.